2001-4596

Onboard Science Analysis and Replanning for Increased Science Return*

M.C. Burl†, S. Chien†, R. Sherwood†, T. Stough†, A. Davies†, P. Zetocha‡

†Jet Propulsion Laboratory,
California Institute of Technology
MS 126-347
Pasadena, CA 91109
{burl,chien}@aig.jpl.nasa.gov

‡Air Force Research Laboratory
AFRL/VSSW, 3550 Aberdeen Ave SE
Kirtland, AFB
Albuquerque, NM 87117-5776
paul.zetocha@kirtland.af.mil

Abstract

Closing the sense-decide-act loop onboard a spacecraft without intervention from the ground will open up tremendous new opportunities in planetary science, space physics, and earth science. Although spacecraft have long been able to make decisions based on self-preservation, e.g., entering safe mode in response to a sensor out-of-bounds, they have lacked the ability to detect and respond to complex patterns and events that are relevant for science purposes. The Autonomous Sciencecraft Constellation (ASC) experiment, which will fly onboard the Air Force's TechSat-21 constellation (an unclassified mission scheduled for launch in 2004), will demonstrate the use of onboard science analysis and replanning to increase science return.

keywords: onboard science, spacecraft autonomy, SAR, change detection, planning, geological feature recognition.

1 Introduction

With current approaches to space exploration, data are downlinked to the ground, analyzed by humans, and then new commands are uplinked to the spacecraft. With this approach, bandwidth constraints and light-time latencies pose significant obstacles to reactive, opportunistic science. As examples of missed scientific opportunities, images from the Voyager missions serendipitously showed volcanic eruptions on Io and new cryo-volcano features on Triton that had not been seen elsewhere in the Solar System, but no follow-up observations of these high-value science targets were possible. Similarly, the Gallileo spacecraft discovered the first asteroid (Ida) having its own moon (Dactyl), but additional measurements that would have been invaluable for more precisely determining the orbit of the asteroidal moon, and hence composition of the main asteroid were not made [17].

NASA autonomy demonstrations, such as the Remote Agent Experiment [21], which executed for several days onboard the NASA Deep Space One mission, and Three Corner Sat (3CS) [8], which is scheduled for launch in 2002, have focused on planning and action, i.e., maintaining the spacecraft in a desired internal state, sequencing lower-level steps together to achieve higher-level goals, and executing plans in a reliable fashion. Although such demonstrations provide valuable milestones in the effort to produce more capable spacecraft, the level of sensory interpretation in these experiments, i.e., the perception phase, has been relatively low. For example, in the RAX experiment, sensing involved checking the state of a switch ("I told the camera to turn on, is it on now?"). However, some of the most compelling arguments for autonomy require that more human-like perceptual capabilities be developed and deployed onboard spacecraft.

The Autonomous Sciencecraft Constellation (ASC) experiment, which is to fly onboard TechSat-21, will flight demonstrate that a remote system can analyze its incoming data streams and autonomously detect and respond to complex spatial objects and change events, which will enable timely onboard decisions and actions leading to greatly increased science return. ASC will also demonstrate the value of an integrated autonomous mission using onboard science analysis, replanning, robust execution, model-based estimation and control, and formation flying. The ASC concept has been selected for flight on the Techsat-21 mission and the necessary software is currently being matured and brought into flight readiness. Key Techsat-21 design reviews occur in Spring 2001 to Spring 2002, with final delivery of the spacecraft and software in September 2003. Nominal launch date is September 2004. The NASA New Millennium Space Technology

6 Project has selected the ASC concept for a Phase-A award. A comprehensive description of the ASC objectives and plan is given in [7]. Here, we will focus more specifically on the onboard science components of ASC.

Section 2 provides an overview of the TechSat-21 mission and the onboard instruments and computing resources that will be available. Each of the three spacecraft will carry an X-band synthetic aperture radar. Section 3 describes the first step in the onboard science analysis, which is to convert the raw, demodulated I and Q returns recorded by the radar into SAR imagery (possibly at reduced resolution). Currently, we are using X-band SAR data of Earth, which was collected by the X-SAR instrument flown with SIR-C on several space shuttle missions in the mid 1990's [15], as a surrogate for actual TechSat-21 data. Three types of science processing algorithms are under consideration for inclusion as part of the flight experiment: recognition of static geological features, change detection, and a discovery/focus of attention algorithm that can identify generically "interesting" regions within a scene. These algorithms are described in Sections 4–6.

Results produced by the science processing algorithms will be used to establish new system-level goals. In early phases of the mission, the range of spacecraft reactions will be limited to more conservative responses such as prioritizing regions around detected features for early downlink. Success in these early tests will enable more aggressive reactions to be evaluated later in the mission, for example, downlinking only data summarizations (e.g., feature catalogs) or data for regions where change has been detected. Although the scientist would never like to discard data, the realities of constrained power and communication bandwidth lead to a situation in which doing so can provide greater science value. Finally, onboard decisions that autonomously alter the planned set of observations in response to high-value science targets, including possibly reconfiguring the constellation, will be considered.

2 Mission Overview

The Air Force Research Laboratory (AFRL) has initiated the TechSat-21 program to serve as a proof of concept for a new class of space missions. The new paradigm seeks to reduce costs and increase system robustness and maintainability by distributing functionality over several micro-satellites flying in formation. The distributed functionality includes processing, command and control, communications, and payload functions. A chief objective is for the system of micro-satellites to in effect function as a "virtual" satellite, which can be controlled and tasked as a single satellite. TechSat-21 is scheduled for a late 2004 launch and will fly three satellites in a near circular orbit at an altitude of 600 Km. The primary mission is one-year in length with the possibility for an extended mission of one or more additional years. During the mission lifetime the cluster of satellites will fly in various configurations with relative separation distances ranging from approximately 100 meters to 5 km. One of the objectives of TechSat-21 is to assess the utility of the space-based, sparse-array aperture formed by the satellite cluster. For TechSat-21, the sparse array will be used to synthesize a large radar antenna. Three modes of radar sensing are planned: synthetic aperture radar (SAR) imaging, moving target indication (MTI), and geo-location.

The principal processor onboard each of the three TechSat-21 spacecraft is a BAE Radiation hardened 175 MIPS, 133MHz PowerPC 750 running the OSE 4.3 operating system from Enea Systems. OSE was chosen because it is inherently message passing based and particularly suitable for distributed applications. Each satellite will have 256 kbytes of EEPROM for boot loads and 128 Mbytes of SDRAM. Communications will be through a Compact PCI bus. For storage of payload data and some large flight software components 8 disk drives per spacecraft will be used. The ASC onboard flight software includes several autonomy software components:

- Onboard science algorithms that will perform SAR image formation, analyze the resulting image data, generate derived science products, and detect trigger conditions such as science events, "interesting" features, and change relative to previous observations.
- Model-based mode identification and execution (MI-R) that uses component-based hardware models to analyze anomalous situations and to generate novel command sequences and repairs.
- Robust execution management software using the Spacecraft Command Language (SCL) [12] package to enable event-driven processing and low-level autonomy .
- The Continuous Activity Planning, Scheduling, and Replanning (CASPER) [6] planner that will replan activities, including downlink, based on science observations in the previous orbit cycles.
- The ObjectAgent and TeamAgent cluster management software will enable the three Techsat-21 spacecraft to autonomously perform maneuvers and high precision formation flying to form

a single virtual instrument.

3 Onboard SAR Image Formation

The image formation module will create a (possibly reduced resolution) SAR image onboard the spacecraft from the raw phase history (demodulated I and Q returns). For the ASC demonstration, we envision only forming a few images per orbit cycle. Although special-purpose real-time SAR processors have been developed (at least for airborne systems) [23], we will carry out the necessary calculations on the general purpose flight processor. Our baseline calculations show that a 10-meter by 10-meter resolution image can be formed in approximately 18 seconds (with full processor utilization). A 2-meter by 2-meter resolution image would require approximately 45 minutes, which still leaves some margin for making decisions within a 90 minute orbit cycle. Note, however, that due to rotation of the Earth between orbits, obtaining imagery of the same region on the ground from consecutive orbit passes, requires that the first orbit use a west-looking imaging geometry and the second pass use an east-looking imaging geometry.

For development and ground-based demonstration purposes, we will use X-SAR data as a surrogate for actual TechSat-21 radar data. The X-SAR instrument was developed under joint sponsorship of the German and Italian governments and was flown, along with the NASA-developed SIR-C (Spaceborne Imaging Radar-C) radar, on several space shuttle missions in the mid-1990's [15]. X-SAR data of several regions of scientific interest including Hawaii, Brazil, and Bangladesh has been obtained from the DLR in both raw and processed forms.

The first step in the onboard science processing is to convert raw data into image form. Although the ground-based processing of the X-SAR data tapes is described in [2], the software was developed for special purpose hardware, so it is not readily available for our purposes. We have proceeded to develop an implementation of the processing based on the description in [2]. Currently, we have implemented only the range compression step, which involves applying a matched filter (linear frequency modulated chirp) to each line of data (refer to discussion in [19]). The chirp parameters are obtained from X-SAR header files [22] to precisely match the characteristics of the transmitted signal. Matched filtering is implemented using the FFT technique, so the complexity of the range compression computation is $O(nm\log(n))$, where m is the number of range lines and n is the number of samples per range line.

A comparison of the image produced with our current processing implementation and the corresponding D-PAF ground-processed product is shown in Figure 1. Note that the onboard image is only single look, while the ground-processed image is an average of multiple looks; hence, the onboard image contains significantly more speckle. Also, the onboard image has not been radiometrically calibrated or focused along track. Nevertheless, the onboard image is already of sufficient quality to identify major features such as the river and other water in the scene. After image formation, several types of science analyses can be performed depending on the current objective.

4 Static Feature Recognition

The ability to autonomously analyze an image to identify specific geological features such as volcanoes, craters, graben, and boulders is needed for a broad set of future missions. For the ASC experiment, a combination of statistical pattern recognition and computer vision techniques will be used to identify features such as volcanoes, lakes, and iceberg fragments.

Rather than hand-coding detectors for each specific object, we approach the problem from a statistical learning point of view. The learning part of the algorithm uses examples provided by a scientist to generate automatically an efficient model for detecting the target object across a range of scales [5] or across a range of variations in appearance [3]. The learning phase will be carried out on the ground using surrogate data (such as X-SAR) or data that will be downlinked from TechSat-21 in the early phases of the mission. The models that are learned can then be uplinked to the spacecraft and used to search for additional instances of the target object.

One type of object model that will be considered for use in ASC is the continuously-scalable detector (CSD) described in [5]. These detectors provide a principled extension of matched filtering to filtering over a continuous scale space. The advantages of matched filtering (or template matching as it is more commonly known in the computer vision and pattern recognition literature) are well known: the procedure is optimal for detecting a known signal in an additive white noise background and the method often works well even when these conditions do not completely hold. CSDs are closely related to the steerable-scalable-deformable filtering ideas developed by Freeman and Adelson [10], Perona [11], and others, as well as to the parametric feature detection work of Nayar, Baker, and Murase [18]. CSDs are ideal for features that are rotationally invariant such as volcanoes and craters. However, the basic method can be extended to handle features with a preferred orien-

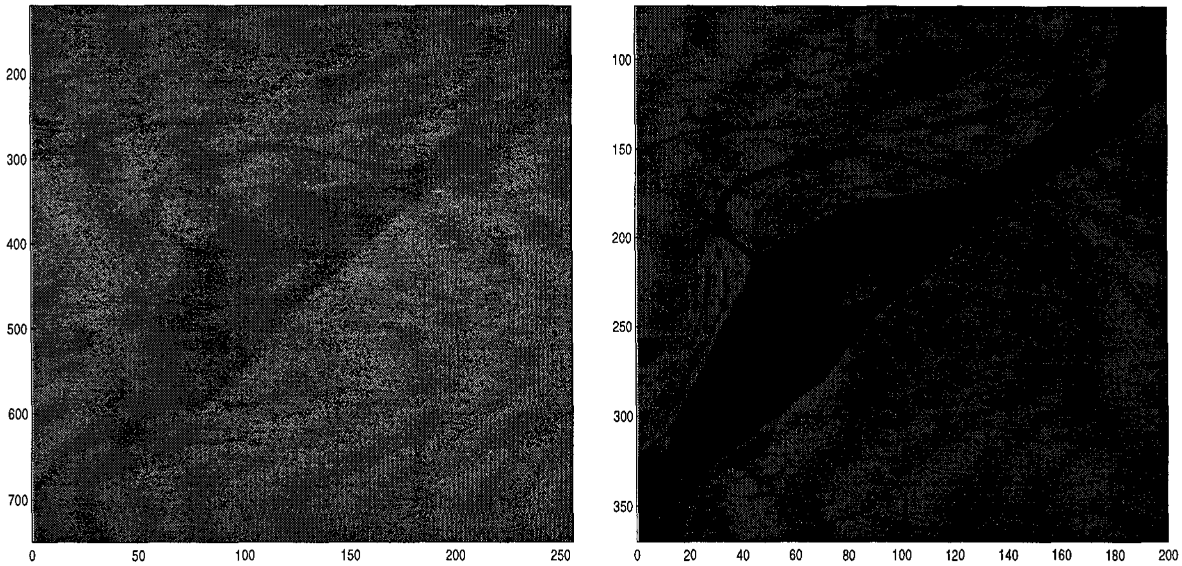

Figure 1: Rio Negro, Brazil: (a) image produced from raw X-SAR data by our current onboard image formation module. (b) corresponding image produced by the DLR D-PAF ground processing. Note that the DLR D-PAF image is an average of multiple looks and hence shows significantly less speckle.

tation such as graben. Merging CSD (or more generally continuously-deformable detectors) with the principal components analysis (PCA) based approach developed in [3] will potentially enable recognition of a more diverse set of objects at various rotations and scales.

Figure 2 illustrates the basic idea behind the CSD, which is to use linear basis functions to approximately represent an object family. The family is formed by resampling a prototype example at different pixel spacings to produce family members at scale 1, 1.01, 1.02, etc. Since there is not much variation from one family member to the next, it is clear that the family has significant redundancy that can be compressed out. The basis functions, which are derived via singular value decomposition (SVD) of the family, provide a compact means to approximately represent the entire family.

5 Change Detection

In addition to calculations based on a single image, the onboard science analysis software will include change detection algorithms that compare images of the same region taken at different times. The change detection capability is particularly relevant for capture of short-term events at the finest time-scale resolutions without overwhelming onboard caching systems and for compressing long-term "monitoring" observations in which changes are infrequent. For space science missions, example applications include tracking atmospheric changes on Jupiter, Neptune, or Triton (from optical image data), tracking ice plate movement on Europa, monitoring known (and identifying new) volcanoes on Io, capturing fine time-scale events such as jet formation on comets or phase transitions in ring systems, and detecting new cratering on planets and moons. To detect change, we will test for statistically significant differences in derived descriptors such as region sizes, locations, boundaries, and histograms, as well as in the raw pixel data. The latter case is complicated by the need to ensure that the two images are precisely co-registered. In part, the orbit repeatability and small absolute positional uncertainty of the TechSat-21 group will help insure co-registration. Since the magnitude of change necessary for the software to declare a trigger event can be specified as a parameter, some degree of robustness to image misalignment will be built in.

We have developed a prototype for detecting change that is based on segmenting approximately-registered "before" and "after" images and comparing gross statistics of the segmented regions. Figure 3 shows a segmentation of an X-SAR image of the Tungi region in Bangladesh into water and non-water. By comparing two such segmentations, for example by monitoring the number of pixels labeled as water in the before and after images, events such as flooding can be detected. Similar methods will be used to segment regions of new lava flow from images taken over Hawaii.

The segmentation algorithm consists of two parts: (1) an unsupervised intensity based clustering step us-

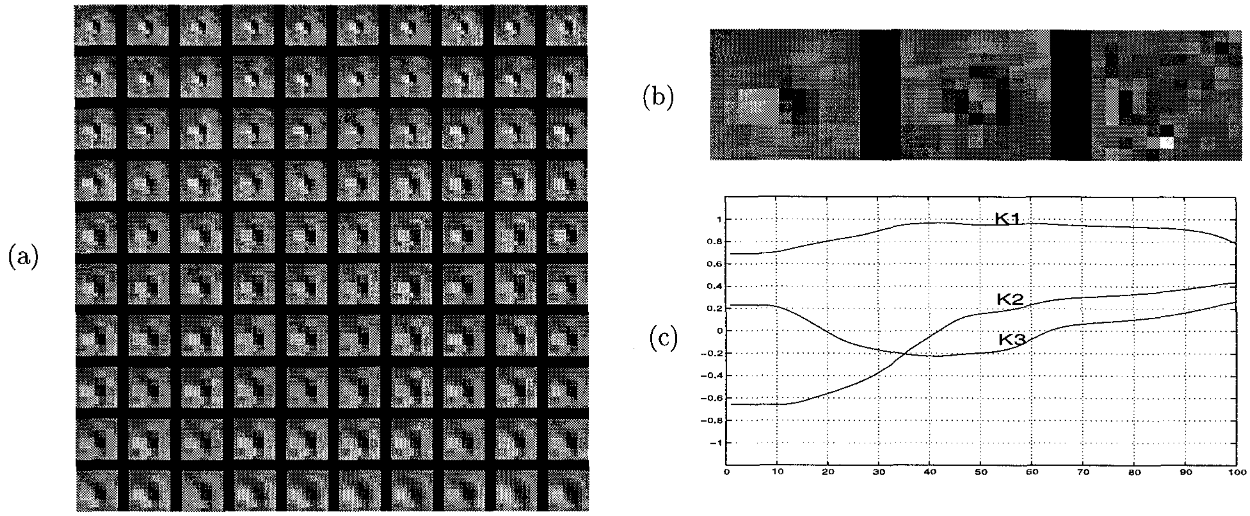

Figure 2: (a) Family of craters constructed by resampling the prototemplate (upper left corner) at finer pixel spacings. (b) First three basis functions obtained through singular value decomposition (SVD) of the crater family. (c) Corresponding interpolating functions as a function of the continuous scale parameter. On the x-axis, 0 corresponds to scale = 1 and 100 corresponds to scale = 2. The y-axis is unitless. The match of an image patch with the template crater at any scale from 1 to 2 can be obtained by computing the inner product of the patch with the three basis functions and combining the results according to the interpolating functions.

ing the well-known k-means clustering algorithm [1, 16] and (2) a Markov Random Field model that takes into account the identity of neighbors to produce a labeling that is spatially more smooth than the pure intensity-based approach. The intensity-based component does not depend on the image being correctly radiometrically calibrated, which can be both an advantage and a disadvantage. If water is present in the scene, then the lowest intensity class identified by the clustering algorithm will correspond to water, but if there is no water in the scene the lowest intensity class may be another type of landcover (e.g., shadow or barren field). With radiometric calibration, the average intensity of the various segments could be compared with statistical data on radar backscatter from water as a function of wavelength and incidence angle to make a more informed decision.

Having change-based triggers will enable key science events to be captured at high temporal (and spatial) resolution without overwhleming the onboard storage capacity. A revolving buffer onboard will cahce several days worth of images of several sites that the spacecraft is monitoring. These images will be collected at the highest spatial and temporal data rates that can be supported by the instrument and solid state recorder. Each day when new images are collected, the oldest images willbe pushed out of the buffer (discarded). However, if the onboard analysis software detects that a major change has occurred, the entire buffer will be latched (preserved) and transferred to a safe area for eventual downlink. In this way, the scientists will be able to see the complete history leading up to the point where change was significant enough to be flagged by the onboard system. Days and days of images where nothing happened will not need to be kept onboard or downlinked, which will lead to a significant increase in the science value that can be returned on a finite-capcity downlink channel.

6 Discovery

Given the increasingly large volumes of data being collected by spacecraft, there is an increasing need for algorithms that can automatically analyze the data streams and focus the spacecraft's attention on "interesting" or salient objects within the data. To complement the work on developing recognizers for specific geological objects, we have also begun to consider generic discovery algorithms [4] that can find objects whose appearance is not known in advance. This capability will be invaluable as spacecraft venture to totally new environments such as Pluto, the surface of Titan, and the subsurface of Europa.

The use of the word "interesting" to describe the type of objects a discovery algorithm should find is a bit loaded since two observers with specific questions and hypotheses in mind may find that completely dif-

(a)

(b)

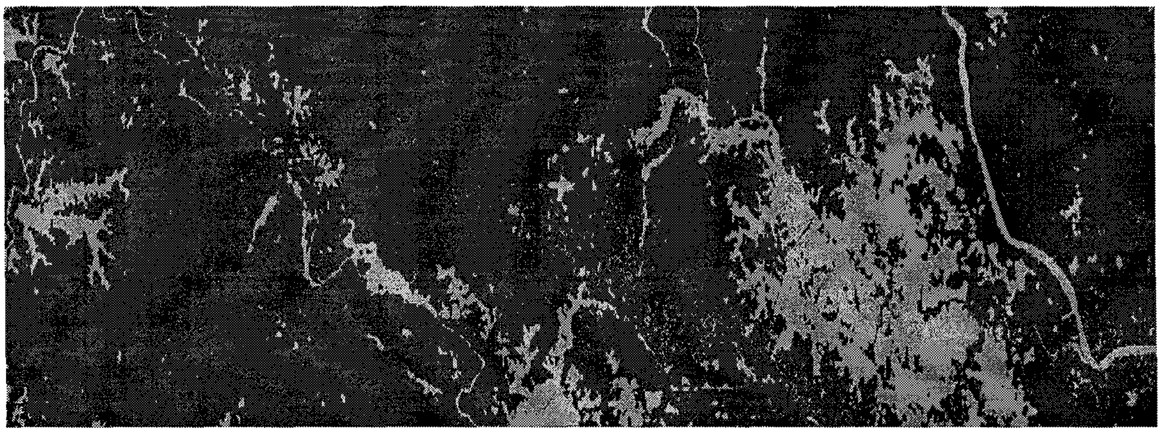

Figure 3: (a) X-SAR image of Tungi region in Bangladesh. (b) Automatic segmentation of the same region. Water regions are associated with the cyan (light-gray in non-color version) labeling. By monitoring counts of pixels labeled as water, ASC will detect changes such as flooding.

ferent portions of an image are interesting. We do not deal this type of top-down, goal-driven interpretation of "interesting". Instead, we take a bottom-up view. It is well known that when people look at an image, their eyes are naturally attracted to certain locations in the image and these locations are fairly consistent across subjects. It is this interpretation of "interesting" that we adopt; thus, our work on visual discovery is closely related to efforts by researchers in biology and psychophysics who are attempting to understand and model the human focus of attention process (e.g., [20, 13, 14] and references contained therein).

The basic approach we have developed is also modeled (somewhat more loosely) after the human visual system. Images are projected into a visually-relevant subspace using a set of multi-orientation, multi-scale Gabor filters that model the receptive field properties of simple cells in the human visual cortex. Within this filter-response subspace, deviant points are identified through an adaptive statistical test that compares the filter-space description at a given spatial location against a model derived from the local background. Deviant points are spatially agglomerated and grouped across scale and orientation. A more detailed description of the discovery algorithm is provided in [4].

The basic discovery algorithm has been tested on a limited, but diverse set of images collected from various solar system bodies via a variety of imaging techniques. In these tests, the algorithm was able to autonomously discover craters on the moon, volcanoes on Venus, sand dunes on Mars, and ice geysers (cryovolcanoes) on Neptune's moon Triton. For all of the tests, the algorithm used the same parameters and was not told what to look for. Instead, it was able to generically identify objects by detecting localized regions whose properties differed significantly from the surrounding background. An example output of the system on a Global Surveyor image of Mars is shown in Figure 4.

In the context of ASC, the discovery capability will be followed by downlink of the identified interesting regions. These regions will enable a scientist to quickly pick out objects within the set that he wants to have re-targeted. Because the algorithm can in some sense direct the scientist to the most interesting regions, human-in-the-loop decisions can be made more quickly and more globally. It will no longer be necessary to wait for every pixel of every image to be downlinked and examined on ground before a human decision can be made. This capability, of course, provides a critical bridge between the current ground-based decisions and the fully autonomous spacecraft that decides on its own which follow-up measurements to make. To be more effective than current practice, the focus of attention mechanism merely needs to return patches that could not be easily obtained by random selection or by a human viewing a highly-compressed version of the entire image.

7 Conclusion

Perceptive spacecraft and robotic explorers will open new avenues of scientific investigation and will significantly impact a broad class of future missions. The ASC flight demonstration of onboard science processing represents a critical first step toward this important goal. ASC will use the flight processor to form SAR images onboard the spacecraft. Then the appropriate (as determined by the onboard planner) science analysis algorithms will be applied to the image(s) to locate specific features, detect changes, or focus attention on interesting areas. The output of the science analysis routines will be used to produce new goals, including requests for follow-up observations. The planner will close the loop onboard by producing a new plan for completing the requested observations. A robust onboard executive will carry out the plans. Autonomously closing the sense-decide-act loop in this way will radically increase science return and open up new opportunities in planetary science, space physics, and earth science.

Acknowledgments

This research has been carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors wish to acknowledge the following TechSat-21/ASC team members: Russell Knight (JPL), Barbara Engelhardt (JPL), Gregg Rabideau (JPL), Ross Wainwright (AFRL), Pete Klupar (AFRL), Pat Cappelaere (Interface and Control Systems), Derek Surka (Princeton Satellite Systems), Brian Williams (MIT), Ronald Greeley (Arizona State University), Victor Baker (U. of Arizona), and James Dohm (U. of Arizona). The authors thank Padhraic Smyth and Igor Cadez for providing the Markov Random Field code used in the segmentation algorithm of Section 5, and Dominic Lucchetti and Charless Fowlkes for their contributions to the discovery algorithm of Section 6. Figure 4 appears with permission from SPIE, copyright 2000.

References

[1] G.H. Ball and D.J. Hall, "A clustering technique for summarizing multivariate data", *Behavioral Science*, **12**, 153–155, March 1967.

[2] R. Bamler, H. Breit, U. Steinbrecher, D. Just, "Algorithms for X-SAR Processing", *IEEE Int. Geoscience and Remote Sensing Digest*, 1993.

[3] M.C. Burl, L. Asker, P. Smyth, U. Fayyad, P. Perona, J. Aubele, and L. Crumpler, "Learning to Recognize Volcanoes on Venus", *Machine Learning Journal*, April 1998.

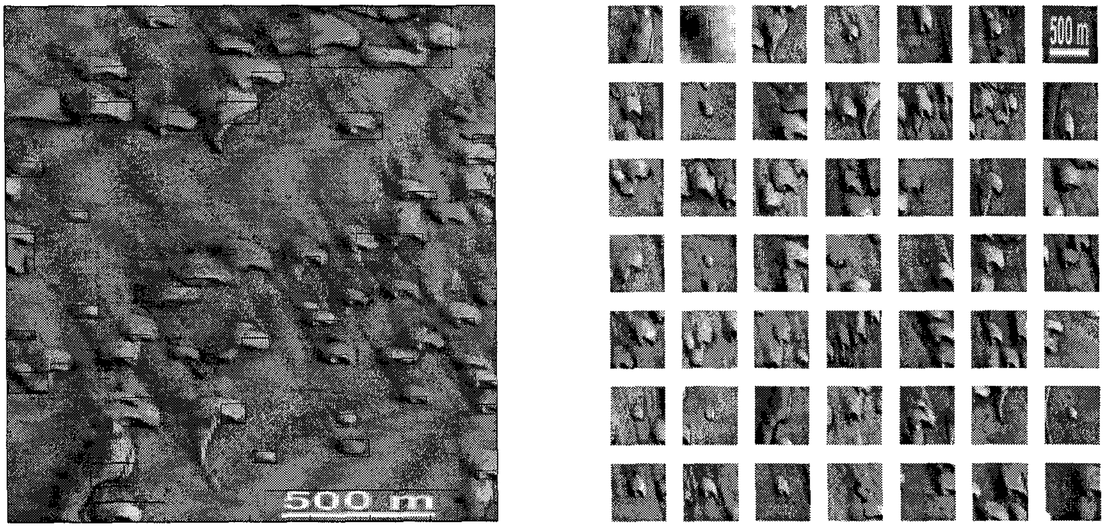

Figure 4: (a) An image of a portion of Mars obtained by Global Surveyor. Many of the identified regions are sand dunes. It is also interesting to note that the "500 m" scale annotation was picked out by the algorithm as interesting. (b) Mosaic of regions of interest.

[4] M.C. Burl and D. Lucchetti, "Autonomous Visual Discovery", *SPIE Aerosense Conf. on Data Mining and Knowledge Discovery*, Orlando, FL, April 2000.

[5] M.C. Burl, W.J. Merline, E.B. Bierhaus, W. Colwell, C.R. Chapman, "Automated Detection of Craters and Other Geological Features", Intl Symp Artificial Intelligence Robotics and Automation in Space, Montreal, Canada, June 2001

[6] S. Chien, G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, G. Stebbins, D. Tran, "ASPEN - Automating Space Mission Operations using Automated Planning and Scheduling", *SpaceOps*, Toulouse, France, June 2000.

[7] S. Chien, R. Sherwood, M.C. Burl, R. Knight, G. Rabideau, B. Engelhardt, A. Davies, P. Zetocha, R. Wainwright, P. Klupar, P. Cappelaere, D. Surka, B. Williams, R. Greeley, V. Baker, J. Dohm, "The Techsat-21 Autonomous Sciencecraft Constellation", *Intl Symp Artificial Intelligence Robotics and Automation in Space*, Montreal, Canada, June 2001.

[8] S. Chien, B. Engelhardt, R. Knight, G. Rabideau, R. Sherwood, E. Hansen, A. Ortiviz, C. Wilklow, S. Wichman "Onboard Autonomy on the Three Corner Sat Mission", *Intl Symposium on Artificial Intelligence Robotics and Automation in Space*, Montreal, Canada, June 2001

[9] C. Elachi, *Intro. to the Physics and Techniques of Remote Sensing*, John Wiley and Sons, 1987.

[10] W. Freeman and E. Adelson, "The design and use of steerable filters", *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 13, pp. 891–906, 1991.

[11] P. Perona, "Deformable kernels for early vision", *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 17, no. 5, pp. 488–499, 1995.

[12] Interface and Control Systems, Spacecraft Command Language, `http://www.sclrules.com`.

[13] L. Itti, C. Koch, and E. Niebur, "A Model of Saliency-Based Visual Attention for Rapid Scene Analysis", *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 20, no. 11, pp. 1254–1259, 1998.

[14] L. Itti and C. Koch, "Learning to detect salient objects in natural scenes using visual attention", *Image Understanding Workshop*, 1999.

[15] R.L. Jordan, B. Huneycutt, and M. Werner, "The SIR-C/X-SAR Synthetic Aperture Radar System", *Proc. of the IEEE*, vol. 79, no. 6, pp. 827–838, 1991.

[16] J. MacQueen, "Some methods for classification and analysis of multivariate observations", in *Proc. Fifth Berkeley Symp. on Math. Stat. and Prob.*, **I**, pp. 281–297, 1967.

[17] W. Merline, *private communication*, 2000.

[18] S. Nayar, S. Baker, and H. Murase, "Parametric Feature Detection", *Comp. Society Conf. on Computer Vision and Pattern Recognition*, pp. 471–477, 1996.

[19] C. Olmsted, "Alaska SAR Facility Scientific SAR User's Guide", Technical Report ASF-SD-003, July 1993.

[20] B.A. Olshausen, C.H. Anderson, and D.C. Van Essen, "A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information", *J. Neuroscience*, vol. 13, no. 11, pp. 4700-4719, Nov. 1993.

[21] Remote Agent Experiment Home Page, `rax.arc.nasa.gov`

[22] "X-SAR CEOS Format", Technical Report produced by the German and Italian PAFs for X-SAR, 1995.

[23] B. Zuerndorfer, G.A. Shaw, "SAR Processing for RASSP Application", *Proc. of First Annual RASSP Conference*, Arlington, VA, Aug 15-18, 1994.